Comments to the National Institute of Standards and Technology (NIST)

RFI Regarding Security Considerations for Artificial Intelligence Agents

March 2026

The Better Identity Coalition appreciates the opportunity to provide comments to the National Institute of Standards and Technology (NIST)) on its RFI Regarding Security Considerations for Artificial Intelligence (AI) Agents.

As background, the Better Identity Coalition is an organization focused on developing and advancing consensus-driven, cross-sector policy solutions that promote the development and adoption of better solutions for identity verification and authentication. Our members – 18 companies in total – are recognized leaders from different sectors of the economy, encompassing firms in financial services, fintech, payments, health care, information technology, and security.

The coalition was launched in February 2018 as an initiative of the Center for Cybersecurity Policy & Law, a non-profit dedicated to promoting education and collaboration with policymakers on policies related to cybersecurity. More on the Coalition is available at https://www.betteridentity.org/.

While our focus as a Coalition has largely been on the "human" side of identity, the rapid advancement of AI agents has posed a number of challenging questions around the ways in which we will manage the identities of agents, as well as how organizations will be able to differentiate agents from humans. Given that compromises of identity systems are the single most exploited attack vector by adversaries in cyberspace – with a recent study noting that identity weaknesses played a material role in almost 90% of investigations into cyber incidents[1] – it is imperative that any discussion on security considerations for AI agents start with identity.

To this point – we believe that figuring out standards-based ways to manage the intersection between human identities and agentic identities is going to be critical to creating a foundation for safe and secure agentic commerce. There are a number of critical challenges here that need to be addressed, including:

- How to enable online service providers to differentiate between a human, an agent that has been authorized by a human to perform a task for them, and an agent or bot that is claiming to have such authorization – but does not? Most companies today invest significant resources in trying to weed out bots from actual human visitors; if agentic commerce is truly going to thrive, those same companies will now need to also sort out "good" bots (such as agents that can be proven to be acting on behalf of a person), vs. "bad" bots that are being used for malicious and/or criminal purposes. Given NIST's longstanding work around digital identity and biometrics, NIST is well-positioned to lead on work that can establish a strong link between an agent and the human who has authorized that agent to act on their behalf.

- How to assign an easily verifiable identity to each agent – in a world where agents might be ephemeral, and only be created to support a specific task for a short period of time? In a world where agents are short-lived and thus have no persistent identifier, how will third parties be able to validate agentic identities? Do existing standards for digital credentials

[1] See https://www.paloaltonetworks.com/resources/research/unit-42-incident-response-report

such as the W3C verifiable credentials standard offer a way to address this challenge, or are new standards needed here?

- How can people delegate authority to an agent to perform some tasks for a person while limiting that agent's ability to do other tasks? Human-to-machine authorization challenges are significant here. With this, there are a number of challenges around identity and authentication to be solved, such as how to enable an agent to authenticate to a site on a person's behalf without giving the agent a person's actual credentials.

- With this, there is also a need to guard against attacks that target the tools used by people to authenticate to agentic systems and authorize actions. As generative AI tools such as deepfakes are increasingly used to target these systems, NIST is well-positioned to lead on standards, guidance, and best practices that can be used to mitigate the risks of these attacks. NIST has already started to address challenges such as liveness detection in its Digital Identity Guidelines, and as deepfake attacks evolve in sophistication, more work will be needed here to ensure that defenders can stay ahead of adversaries. In addition, the Financial Services Sector Coordinating Committee (FSSCC) recently published guidance on this topic that may be helpful.[2]

- If something goes wrong in an agentic commerce use case – such as an agent using a compromised identity in a transaction, or going beyond the parameters of the agent's authorization (such as ordering 4,000 pounds of meat when the agent was asked to find cheap hamburgers for lunch) – what tools will enable a party to figure out what went wrong and who should be held responsible?

Given the myriad challenges we have faced in trying to solve digital identity issues for humans, we have concerns that a world where we are unleashing millions of AI agents that will have their own identities – and that will be insisting to the companies and other organizations we do business with that they are authorized to be acting our behalf – is likely to be fraught with problems if a foundation of standards and best practices to manage agentic identity does not exist.

Here, we are greatly encouraged by two recent announcements from NIST:

1) First, we are excited about NIST's recent announcement that it is considering launching a new project at the National Cybersecurity Center of Excellence (NCCoE) focused on Software and AI Agent Identity and Authorization.[3] From our perspective, moving forward with this NCCoE project is the single most important step that NIST (and the Federal government) can take to help address challenges like the ones we outlined above.

[2] See Mitigating AI-Powered Attacks Against Identity and Authentication" at https://fsscc.org/wp-content/uploads/2026/02/AI-IA-Workstream-Mitigations.pdf
[3] See https://www.nccoe.nist.gov/projects/software-and-ai-agent-identity-and-authorization

Our primary feedback on NIST's work here is that it does not go deep enough. In that the current NCCoE project anticipates tackling only enterprise use cases, whereas there is another, complementary set of use cases involving consumer-facing use cases of agents where standards and best practices are also desperately needed.

While we strongly support NIST's initial focus on enterprise uses cases, we would urge NIST to not only fund and expeditiously advance the proposed NCCoE project, but also either expand its scope to handle consumer use cases or launch a second, complementary, NCCoE project to focus specifically on consumer use cases.

2) Second, we were also pleased to see the recent launch of the "AI Agent Standards Initiative" out of the Center for AI Standards and Innovation (CAISI) at NIST.

   As noted above, we believe that figuring out standards-based ways to manage the intersection between human identities and agentic identities is going to be critical to creating a foundation for safe and secure agentic commerce. While the ultimate goal here should be the development of voluntary, consensus-based standards whose development is led by bodes outside of government, NIST has a critical role to play here in helping to facilitate the industry-led development of standards and ensuring that U.S. interests are represented in international standards bodies.

   The good news here is that standards efforts are underway to tackle agentic identity issues in a variety of standards bodies, providing NIST with ample opportunities to engage and lead. However, these opportunities will be lost if NIST's standards engagement work on agentic identity is not properly funded and staffed, and we believe more resources are needed here.

We greatly appreciate NIST's willingness to consider our comments and suggestions, and welcome the opportunity to have further discussions. Should you have any questions on our feedback, please contact the Better Identity Coalition's coordinator, Jeremy Grant, at jeremy.grant@venable.com.